Predicting Local-Official Corruption by Coupling Rational Choice Theory with XGBoost: A Computational Sociology Approach

Fajerin Biabdillah, Arsan Kumala Jaya *

Politeknik Negeri Samarinda, Samarinda, East Kalimantan, Indonesia
fajerin@polnes.ac.id, arsankumalajaya@polnes.ac.id
* Corresponding author: arsankumalajaya@polnes.ac.id

Abstract. Risk-based audit prioritisation requires a model that ranks officials before an investigation begins, not one that recognises officials already charged. This paper reformulates the prediction of local-official corruption in Indonesia as a strictly prospective task and evaluates it under leakage controls appropriate to that claim. We construct the Indonesian Corruption Corpus, an official-year panel of 545 governors, regents and mayors observed from 2015 to 2024, linking asset filings, budget realisations, electoral family records and adverse legal outcomes. Rational Choice Theory is translated into four observable proxies, including a Dynastic Insulation Index (DII) and Wealth Velocity. Features observed up to year t predict a first adverse legal outcome in t+1 or t+2; official-years at or after that outcome are censored, all preprocessing is fitted inside training folds, and officials are disjoint across training and hold-out windows. XGBoost attains a hold-out ROC-AUC of 0.71 (95% CI 0.63–0.78) and a PR-AUC of 0.18 against a 2.78% base rate, recovering 40.4% of prospective cases at a 10% audit capacity — a fourfold gain over random selection. The DII result survives five alternative specifications, but officials with any dynastic tie face roughly twice the false-positive rate. The system is defensible only as a contestable, human-reviewed prioritisation aid.

Keywords: Computational Sociology, Corruption Risk Prediction, Gradient Boosting, Rational Choice Theory, Political Dynasties, Algorithmic Accountability.

1 Introduction

Decentralisation after 1998 moved substantial fiscal authority to Indonesian district and provincial executives, and a body of work documents that it redistributed opportunities for rent extraction with it [1, 2]. Anti-corruption capacity has not scaled likewise: the audit body and the anti-corruption commission allocate a limited inspection workforce largely through complaint-driven and rotational selection [3, 4], so the selection rule governs how much malfeasance is observed at least as strongly as the underlying rate. This motivates risk scoring, evaluated elsewhere on procurement contracts in Italy [5], Colombia [6] and across national administrations [7, 8], on aggregate corruption indices [9], on Spanish provinces [10] and on petty-corruption intention among Indonesian officers [11]. Two gaps remain. Existing systems operate on the contract or the jurisdiction, while the individual officeholder — the unit an auditor schedules — is rarely modelled directly. And many published classifiers label an entity positive whenever a case is known, including cases already public when the predictors were recorded, so the model can score highly while learning the signature of an investigation rather than the conditions preceding one.

We therefore enforce the temporal ordering: predictors are restricted to information available by the end of year t, the outcome is a first adverse legal event in t+1 or t+2, every official-year at or after that event leaves the sample, preprocessing is fitted within training folds, and officials are not shared between training and hold-out windows. Performance falls well below the retrospective figures reported previously, which we regard as the

© The Author(s) 2026
M. Murtadha et al. (eds.), *Proceedings of the 2026 11th International Conference on Sustainable Information Engineering and Technology (SIET 2026)*, Atlantis Highlights in Computer Sciences 26,
https://doi.org/10.2991/978-94-6239-828-3_31

consequence of removing a confound rather than a loss. We contribute a construct-to-proxy-to-feature mapping (Sect. 2), a documented corpus and entity-resolution protocol with a data dictionary (Sect. 3), a prospective leakage-controlled evaluation (Sects. 4–5), and a robustness study of the Dynastic Insulation Index with a subgroup analysis of who bears the false positives. Four questions follow. RQ1: can features up to t predict a first adverse outcome in t+1 or t+2 usefully under censoring, grouped validation and official-level disjointness? RQ2: do the DII and Wealth Velocity contribute beyond conventional indicators, and does that survive alternative index specifications? RQ3: are outputs calibrated enough for probability-based tiers, and what recall is attainable at staffable capacities? RQ4: how is the false-positive burden distributed across office type, region and dynastic status?

2 Theoretical Framework

Following Becker's economic model of offending [12] and Klitgaard's decomposition of corruption into monopoly, discretion and accountability [13], we treat an officeholder as choosing between a compliant and a non-compliant course, writing E[U_c] for the expected utility of the latter:

$$E[U_c] = (1 - P[det]) \cdot Y[corr] - P[det] \cdot (C[pen] + C[soc]) \quad (1)$$

In (1), P[det] is the probability that conduct is detected and pursued, Y[corr] the material gain conditional on non-detection, C[pen] the statutory sanction and C[soc] the reputational cost. These symbols are used unchanged throughout and each maps to one proxy family in Table 1; work on Indonesian local government suggests P[det] and C[soc] are both attenuated where political families are entrenched [2, 14]. We keep three levels distinct, because conflating them is a common weakness here. A construct is a quantity in (1) and is never observed; a proxy is an administrative quantity argued to co-vary with it; a feature is the column supplied to the classifier, and none enters without an entry in Table 1. A proxy is not a measurement of its construct: a high Cash-to-Asset Ratio is as consistent with liquidity preference or inheritance as with concealment.

3 The Indonesian Corruption Corpus

3.1 Construction and Entity Resolution

The frame is every governor, regent and mayor holding office for at least one full calendar year between 2015 and 2024, requiring an identifiable term in Ministry of Home Affairs records, an LHKPN filing during or within a year of the term, and a resolvable KPU candidacy record; acting officeholders serving under six months were excluded. Of 618 officeholders, 51 lacked LHKPN records, 14 had unresolvable candidacy records and 8 were short-term appointments, leaving 545 officials and 4,360 official-year rows. Linkage ran deterministically on national identity number and then on the exact triple (normalised name, date of birth, jurisdiction code), followed by Jaro-Winkler matching blocked within jurisdiction and term, accepting above 0.94 and adjudicating 0.88–0.94 manually. Of 312 adjudicated pairs 47 were rejected, with Cohen's kappa of 0.91 on a 150-pair double-coded subsample.

3.2 Outcome Variable and Effective Sample Size

An adverse legal outcome is recorded when an official is named a suspect (tersangka) by the anti-corruption commission or the Attorney General, or convicted under Law No. 31/1999, dated at the earlier event. Because the design is prospective only the first such event per official is used; 64 of the 545 officials experience one. After censoring and the one-year lag needed for trajectory features the modelling sample is 3,914 official-years from 541 officials, of which 113 (2.89%) are prospective

positives. Those 113 rows carry only 64 independent events, so all intervals in Sect. 5 are clustered on the official.

3.3 Features and Their Limits

Four proxies are theory-derived: Wealth Velocity (3), the Dynastic Insulation Index (4), the Cash-to-Asset Ratio (liquid holdings over declared net worth) and a Benford deviation — the chi-squared distance between the first-digit distribution of procurement awards issued under an official's authority in year t and the Benford reference. Following standard forensic-accounting practice it is computed only where at least 200 awards are recorded in that official-year and set missing otherwise rather than imputed; about 22% of official-years fall below the threshold. Missingness is not noise: each proxy family carries a companion binary indicator. Because an LHKPN filing for year t is normally published the following year, a filing is admitted as a predictor for t only if its publication date precedes the outcome window, and the same rule applies to audited APBD realisations.

Table 1. Data dictionary: construct, feature, source and year of availability.

Construct	Feature (unit)	Source	Avail.
P[det]	DII (0–1); relatives in office (count)	KPU	t
Y[corr]	Wealth Velocity (ratio); Cash-to-Asset (share)	LHKPN	t (pub. t+1)
Y[corr]	Real net-worth change (IDR bn, 2018)	LHKPN, BPS	t (pub. t+1)
Discretion	Grant/social-aid share; election-year spike	APBD, KPU	t (aud. t+1)
Discretion	Benford deviation (chi-squared)	LPSE	t
C[soc]	Term number (count); win margin (share)	KPU	t
Context	Log GRDP per capita; missingness flags (x4)	BPS, derived	t

4 Prospective Design and Leakage Control

4.1 Label Definition and Sample Partition

Let T[i] be the year of official i's first adverse legal outcome, or infinity where none occurs. For each official-year (i, t):

$$Y[i,t] = 1 \text{ if } T[i] \in \{ t+1, t+2 \}; \quad Y[i,t] = 0 \text{ if } T[i] > t+2; \quad (i,t) \text{ censored if } t \geq T[i] \quad (2)$$

The censoring clause in (2) separates this design from the retrospective formulation: once T[i] has occurred no further row for that official enters the sample, so the classifier never sees an official-year whose covariates were recorded after the case became public. Training uses feature years 2015–2019 (outcomes through 2021) and hold-out 2020–2022 (outcomes through 2024). Officials appearing in both windows are assigned exclusively to training, so none contributes to both: training is 2,043 official-years (61 positives, 35 officials) and hold-out 1,871 (52 positives, 29 officials), which therefore measures generalisation to unseen officials in unseen years.

4.2 Leakage-Controlled Fitting

Cross-validation uses five grouped folds keyed on official identity, required because official-years for one person are strongly serially correlated and a random split would place adjacent years of the same official on both sides [15]. Every data-dependent operation is fitted on the training partition and applied unchanged to the validation partition: imputation, deflation and scaling, feature construction, and the decision threshold. Hyperparameters are selected in an inner loop nested within each outer fold, and isotonic calibration is fitted on inner-fold predictions only. A missing W[i,t] is replaced by W[i,t−1] scaled by one plus the mean year-on-year growth of officials

in the same regional cluster; that peer set is restricted to the training partition and to training-window years, so no imputed value borrows from a validation or hold-out year.

4.3 Feature Definitions

Wealth Velocity expresses the change in real declared net worth as a multiple of statutory annual income:

$$V[i,t] = (W[i,t] - W[i,t-1]) / \max(I[off,i,t], \varepsilon) \quad (3)$$

Both W and I[off] in (3) are in constant 2018 rupiah, so V[i,t] is dimensionless — multiples of annual official income — with $\varepsilon = 1$ IDR guarding the denominator. V is never reported in currency units; where a monetary quantity is needed, Sect. 5.5 reports the real net-worth change separately in IDR billion, correcting an inconsistency in the previous version. The Dynastic Insulation Index aggregates concurrent officeholding among relatives; with R[i] the relatives of i, p(j) the position held by relative j and w(·) a position weight,

$$DII[i,t] = \min \{ 1 , \Sigma \text{ over } j \in R[i] \text{ of } w(p(j)) \cdot 1[p(j) \neq 0 \wedge jur(j) = jur(i)] \} \quad (4)$$

The indicator in (4) restricts the sum to relatives holding office in an overlapping jurisdiction, and w(·) weights positions by influence over local enforcement and budget approval (national legislature 1.0, governor 0.9, regent or mayor 0.8, local legislature 0.5, agency head 0.4, state-enterprise head 0.3). These weights are ex ante and partly normative, which is why Table 4 re-estimates the model under alternative specifications and Sect. 6.3 treats the index as the paper's most contestable component.

4.4 Model and Baselines

The classifier is XGBoost [16], minimising a regularised binary log-loss by stagewise addition of trees, with scale_pos_weight set to the within-fold negative-to-positive ratio. The grid searched learning rate {0.01, 0.05, 0.1}, depth {3, 5, 7}, minimum child weight {1, 3, 5}, subsample {0.7, 0.85, 1.0}, column subsample {0.6, 0.8, 1.0} and gamma {0, 0.5, 1.0} over 500–1500 rounds with early stopping; the modal selection was 0.05, depth 3, child weight 3, subsample 0.85, column subsample 0.8, gamma 0.5. Baselines are L2 logistic regression, a 500-tree random forest and LightGBM [17], fitted under identical folds and preprocessing. Fig. 1 summarises the complete pipeline and marks the barrier beyond which nothing is fitted outside a training fold.

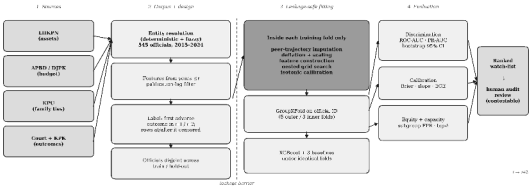


Fig. 1. Prospective risk pipeline. Features come from years up to t, outcomes from t+1 to t+2. Everything right of the dashed barrier is fitted inside training folds only.

5 Results

5.1 Discrimination and Calibration

Table 2 reports grouped cross-validation and the temporal hold-out with 95% intervals from 2,000 bootstrap resamples clustered on the official, the appropriate unit given that 52 hold-out positives arise from 29 officials. XGBoost attains a hold-out ROC-AUC of 0.71 [.63–.78] and PR-AUC of 0.18 [.11–.27] against a 2.78% base

rate, a 6.5-fold improvement in average precision over the no-skill value; PR-AUC is reported because ROC-AUC is optimistic under this imbalance [18]. The XGBoost interval overlaps LightGBM's, so we claim no library-level advantage. Calibration matters because Sect. 5.5 tiers by probability: uncalibrated scores are over-dispersed (slope 0.71, expected calibration error 0.041), while isotonic regression on inner folds gives slope 0.98 and error 0.013, improving the Brier score from 0.0261 to 0.0247.

Table 2. Prospective performance. CV: mean ± SD over five grouped folds. Hold-out: estimate with 95% cluster-bootstrap interval. Base rate 2.78%.

Model	AUC (CV)	PR (CV)	AUC hold-out	PR hold-out
Logistic	0.63 ± .04	0.11 ± .03	0.61 [.53–.69]	0.09 [.05–.15]
Random forest	0.70 ± .03	0.16 ± .03	0.68 [.60–.75]	0.14 [.08–.22]
LightGBM	0.73 ± .03	0.20 ± .04	0.70 [.62–.77]	0.17 [.10–.26]
XGBoost	**0.74 ± .03**	**0.21 ± .04**	**0.71 [.63–.78]**	**0.18 [.11–.27]**

5.2 Operating Points at Realistic Audit Capacity

A ranking model is used at a capacity an agency can staff, so Table 3 reports confusion counts at three top-k budgets rather than at the default 0.5 cut-off, which is uninformative at this base rate. At 10% capacity the model flags 187 of 1,871 hold-out official-years, recovers 21 of 52 prospective cases (40.4%) at 90.9% specificity, and gives a 4.0-fold gain over random selection. The 11.2% precision means roughly nine reviews per confirmed case — the quantity that determines deployability.

Table 3. Hold-out operating points by audit capacity (1,871 rows, 52 positives).

Cap.	Flag	TP	FP	FN	Recall	Prec.	Spec.	Lift
5%	94	13	81	39	25.0%	13.8%	95.5%	5.0x
10%	**187**	**21**	**166**	**31**	**40.4%**	**11.2%**	**90.9%**	**4.0x**
20%	374	31	343	21	59.6%	8.3%	81.1%	3.0x

5.3 Ablation and Robustness of the Dynastic Insulation Index

Table 4 removes one feature family at a time and refits the entire pipeline, hyperparameter selection included, inside the same folds. The theory-derived features remain load-bearing: removing either costs 0.03 ROC-AUC and removing both costs 0.06, answering RQ2 affirmatively. Two smaller ablations are reported rather than tabulated: dropping the Benford and budget-anomaly block costs 0.01, and dropping the missingness indicators costs 0.02, so non-filing is informative rather than a nuisance to impute away. Because the weights in (4) are partly normative, the lower block re-estimates the model under alternative constructions. Hold-out ROC-AUC ranges from 0.69 to 0.72 against the reference 0.71 (equal-weight and province-normalised variants, omitted for space, give 0.70 and 0.71), every variant keeping a clear margin over removing the index entirely (0.68). Differences among graded variants sit inside the bootstrap interval: the result is driven by the presence of dynastic ties rather than the particular weights.

Table 4. Ablation (CV) and Dynastic Insulation Index sensitivity (hold-out); each variant refits preprocessing and hyperparameters in-fold.

Variant	AUC	PR-AUC	Δ AUC
Ablation: full feature set (CV)	**0.74**	**0.21**	—
− Dynastic Insulation Index	0.71	0.17	−0.03
− Wealth Velocity	0.71	0.18	−0.03
− DII and Wealth Velocity	0.68	0.14	−0.06
DII spec.: ex-ante weights (hold-out)	**0.71**	**0.18**	—
Unweighted count of relatives	0.70	0.17	−0.01
Binary dynasty indicator	0.69	0.15	−0.02
In-fold fitted weights	0.72	0.19	+0.01

5.4 Who Bears the False Positives

Because the model conditions on family networks, it can concentrate review on officials from politically connected families irrespective of conduct. Table 5 disaggregates the 10% operating point. Officials with any dynastic tie face a 14.2% false-positive rate against 6.8% for those with none, a ratio of 2.09; regional rates are closer (1.28) and office-type differences modest (governors 1.21). Group-specific thresholds equalising false-positive rates cut the dynastic ratio to 1.18 at a cost of 4.6 percentage points of recall. We report both configurations rather than the more favourable one alone; choosing between them is a policy decision about acceptable burden, not a modelling optimum [19].

Table 5. Subgroup performance at 10% capacity, hold-out. Ratio is FPR against each reference level.

Subgroup	Share	Recall	FPR	Ratio
Dynastic tie present (DII > 0)	31.4%	48.1%	14.2%	2.09
No dynastic tie (ref.)	68.6%	32.0%	6.8%	1.00
Java–Bali	38.2%	42.3%	10.4%	1.28
Outside Java–Bali (ref.)	61.8%	39.1%	8.1%	1.00
Equalised-FPR thresholds	**100%**	**35.8%**	**9.1%**	**1.18**

5.5 Calibrated Risk Tiers

The watch-list groups official-years into four calibrated bands whose observed rates track them closely: Critical ($P > 0.30$) covers 2.1% of rows at a 24.4% observed rate, High ($0.15 < P \leq 0.30$) 7.9% at 9.7%, Moderate ($0.05 < P \leq 0.15$) 22.6% at 3.4%, and Low ($P \leq 0.05$) the remaining 67.4% at 1.0%. This correspondence licenses reporting the bands as probabilities rather than bare scores. Mean Wealth Velocity falls monotonically across them (6.8, 3.1, 1.2, 0.4) as the dimensionless ratio of (3), while the corresponding median real net-worth changes are 3.9, 1.4, 0.4 and 0.1 IDR billion — a separate monetary quantity.

6 Discussion

6.1 The Model Predicts Detection, Not Corruption

The outcome records enforcement activity, so the classifier estimates the probability that an official-year is followed by an adverse legal event, not that misconduct occurred. Officials never investigated are labelled negative whether or not they offended — a positive-and-unlabelled problem rather than a clean binary one [20]. Reported precision is therefore a lower bound on genuine concern among flagged officials, and the model inherits whatever selectivity the agencies exhibit, so Table 5 is partly a measurement of that selectivity. Feature attributions likewise describe the model, not the world: a large attribution for the DII means the index improves predictions of enforcement outcomes, not that family ties cause corruption.

6.2 Contestability

A flagged official-year should enter a defined administrative process, not a conclusion. Every flag should carry the feature values that produced it; the official should be notified once a review opens and be able to submit correcting documentation, particularly for asset and family records where source errors are common; and reviews closing without finding should be recorded, so recurrent unproductive flags become visible in audit. Without this the disparity in Table 5 becomes a standing burden on an identifiable group, the characteristic failure mode of public-sector risk scoring [19].

6.3 Limitations

The Dynastic Insulation Index is the most contestable element of the framework: its weights are chosen ex ante rather than estimated, its coverage depends on KPU filings in which an unrecorded relative is indistinguishable from an absent one, and it treats jurisdiction as an administrative boundary when influence may run through informal networks. Table 4 shows the predictive conclusion is stable across specifications, but stability is not validity: no variant establishes that the index measures political insulation rather than visibility to investigators. Three further limits apply: the effective sample is 64 events, which is why intervals are wide; the hold-out spans a single institutional period; and external validity beyond Indonesian local executives is untested, making cross-regional replication [5, 8] the next step.

6.4 Reproducibility and Data Availability

LHKPN filings, APBD realisations, KPU records and court decisions are individually public, but the linked corpus identifies named individuals in relation to legal proceedings and is not released in that form. We make available the full feature schema extending Table 1; the entity-resolution, feature-construction and evaluation code; a synthetic corpus of 545 simulated officials reproducing the marginal distributions, missingness pattern and class balance of the ICC, on which the pipeline runs end to end; and aggregated fold-level outputs underlying Tables 2–5. Supervised access to the identified corpus is handled by the corresponding author under a data-use agreement.

7 Conclusion

Reformulating local-official corruption prediction as a prospective task lowers measured performance and raises what can honestly be claimed from it. Under censoring, grouped validation, official-level disjointness and strictly in-fold preprocessing, XGBoost reaches a hold-out ROC-AUC of 0.71 [.63–.78] and PR-AUC of 0.18 against a 2.78% base rate, recovering 40.4% of prospective cases at 10% audit capacity — a fourfold gain over random selection (RQ1, RQ3). The two Rational-Choice-derived features carry the predictive weight and survive five alternative index specifications (RQ2). The false-positive burden falls about twice as heavily on officials with dynastic ties, a disparity group-specific thresholds reduce but do not remove (RQ4). The model estimates the risk of enforcement, not of wrongdoing, and is defensible only as a contestable input to human audit scheduling.

Acknowledgments. The authors thank colleagues at Politeknik Negeri Samarinda, and the reviewers, whose objections to the label design and to the treatment of dynastic features shaped this revision.

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Aulia, A.A.: Disruption in corruption eradication in Indonesia. Public Integr. 27, 736–757 (2025). doi:10.1080/10999922.2025.2455757
2. Silitonga, M.S., Wittek, R., Snijders, T.A.B., Heyse, L.: Democratizing corruption: a role structure analysis of Indonesia's Big Bang decentralization. Appl. Netw. Sci. 8 (2023). doi:10.1007/s41109-023-00535-w
3. Prihatmanto, H.N., Munajat, M.D.E., Irawati, I.: Detecting the corruption pattern and measuring the corruption detection pace at the Indonesian village level. J. Tata Kelola Akunt. Keuang. Negara 9, 289–308 (2023). doi:10.28986/jtaken.v9i2.1361
4. Paranata, A.: A systematic literature review of anti-corruption policy: a future research agenda in Indonesia. Public Organ. Rev. 25, 1181–1214 (2025). doi:10.1007/s11115-025-00847-8
5. Decarolis, F., Giorgiantonio, C.: Corruption red flags in public procurement: new evidence from Italian calls for tenders. EPJ Data Sci. 11, 16 (2022). doi:10.1140/epjds/s13688-022-00325-x

6. Gallego, J., Rivero, G., Martínez, J.D.S.: Preventing rather than punishing: an early warning model of malfeasance in public procurement. Int. J. Forecast. 37, 360–377 (2020). doi:10.1016/j.ijforecast.2020.06.006
7. Fazekas, M., Kocsis, G.: Uncovering high-level corruption: cross-national objective corruption risk indicators using public procurement data. Br. J. Polit. Sci. 50, 155–164 (2020). doi:10.1017/S0007123417000461
8. Ash, E., Galletta, S., Giommoni, T.: A machine learning approach to analyze and support anticorruption policy. Am. Econ. J. Econ. Policy 17, 162–193 (2025). doi:10.1257/pol.20210618
9. Lima, M.S.M., Delen, D.: Predicting and explaining corruption across countries: a machine learning approach. Gov. Inf. Q. 37, 101407 (2020). doi:10.1016/j.giq.2019.101407
10. López-Iturriaga, F.J., Pastor Sanz, I.: Predicting public corruption with neural networks: an analysis of Spanish provinces. Soc. Indic. Res. 140, 975–998 (2018). doi:10.1007/s11205-017-1802-2
11. Masrom, S., Rahman, R.A., Salleh, N.A., et al.: Machine learning prediction of petty corruption intention among law enforcement officers. Indones. J. Electr. Eng. Comput. Sci. 30, 1634–1642 (2023). doi:10.11591/ijeecs.v30.i3.pp1634-1642
12. Becker, G.S.: Crime and punishment: an economic approach. J. Polit. Econ. 76, 169–217 (1968). doi:10.1086/259394
13. Klitgaard, R.: Controlling Corruption. University of California Press, Berkeley (2019). doi:10.1525/9780520911185
14. Olken, B., Pande, R.: Corruption in developing countries. Annu. Rev. Econ. 4, 479–509 (2012). doi:10.1146/annurev-economics-080511-110917
15. Kaufman, S., Rosset, S., Perlich, C., Stitelman, O.: Leakage in data mining. ACM Trans. Knowl. Discov. Data 6, 1–21 (2012). doi:10.1145/2382577.2382579
16. Chen, T., Guestrin, C.: XGBoost: a scalable tree boosting system. In: KDD '16, pp. 785–794. ACM (2016). doi:10.1145/2939672.2939785
17. Ke, G., Meng, Q., Finley, T., et al.: LightGBM: a highly efficient gradient boosting decision tree. In: NeurIPS 30, pp. 3146–3154 (2017).
18. Saito, T., Rehmsmeier, M.: The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. PLoS ONE 10, e0118432 (2015). doi:10.1371/journal.pone.0118432
19. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. In: NeurIPS 29, pp. 3315–3323 (2016).
20. Elkan, C., Noto, K.: Learning classifiers from only positive and unlabeled data. In: KDD '08, pp. 213–220. ACM (2008). doi:10.1145/1401890.1401920

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

